

**DATA MANAGEMENT SYSTEMS, DATA MANAGEMENT SYSTEM STORAGE
DEVICES, ARTICLES OF MANUFACTURE, AND
DATA MANAGEMENT METHODS**

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is related to copending U.S. Patent Application entitled "Data Management Systems, Articles Of Manufacture, And Data Storage Methods" having Attorney Docket No. 100204110-1, listing Stephen Gold, Harald Burose, and Sebastien Schikora as inventors, which was filed concurrently and which is incorporated by reference herein.

EL979977656
FIELD OF THE INVENTION

[0002] At least some embodiments of the invention relate to data management systems, data management system storage devices, articles of manufacture, and data management methods.

BACKGROUND OF THE INVENTION

[0003] Reliable data protection solutions are becoming increasingly important as the usage of and reliance upon electronic data continue to increase. For example, data of numerous business processes may be entirely or substantially electronically stored. Exemplary processes including accounting systems, administrative (e.g., docketing, calendaring, recordkeeping, etc.) systems, email systems, etc. may generate significant amounts of electronic data. It is important to protect against loss of the generated electronic data.

[0004] Accordingly, high capacity data protection storage devices are becoming commonplace in numerous implementations, such as libraries, and backup and archive configurations, for example. It is desired for these arrangements to provide storage capabilities with minimal interrupts of the host system being protected. Increased speed, accessibility and reliability are additionally desired for providing short or long-term storage of digital data.

[0005] Various types of storage media may be used in different data protection solutions. Exemplary types of media include disk and/or tape, although others may be used. Tape systems provide benefits of relatively

simplistic design and reduced cost, but suffer from drawbacks of having relatively inflexible designs, and relatively slow speeds.

[0006] During an initial backup operation, data protection storage systems may store baseline versions of files. Thereafter, delta versions of the files may be created during subsequent backup operations which define differences between the subsequent versions of the files and the original versions of the respective files. If the capacity of a subsystem of a data protection storage system is consumed, backups with respect to the subsystem are halted, or data is moved to another subsystem and the changes may be recorded for subsequent retrieval. Such movement of data may be significant (e.g., all baselines and deltas for an entire protected computer dataset). Movement of significant amounts of data may require a relatively large amount of data traffic between plural subsystems and have a major impact on performance of the data protection storage system.

SUMMARY OF THE INVENTION

[0007] At least some embodiments of the present invention provide data management systems, data management system storage devices, articles of manufacture, and data management methods.

[0008] According to one embodiment of the invention, a data management system comprises a plurality of storage devices individually comprising a physical storage space, wherein the physical storage space of one of the storage devices is configured to store a baseline version of a data object and the physical storage space of another of the storage devices is configured to store a delta version of the data object and processing circuitry configured to control storage operations of at least one of the storage devices, to process a restore request with respect to the data object, to access the delta version from the other of the storage devices responsive to the restore request, and to initiate communication of data of the baseline version and the delta version of the data object to a computer system.

[0009] According to another embodiment of the invention, an article of manufacture comprises a processor-usable medium comprising processor-usable code configured to cause processing circuitry of one of a plurality of storage

devices of a data management system to receive a request to store a baseline version of a data object, effect storage of the baseline version using physical storage space of the one of the storage devices, receive a request to store a delta version of the data object after the effecting storage of the baseline version, accessing information regarding a status of the one of the storage devices, and initiate storage of the delta version using an other of the storage devices of the data.

[0010] According to an additional embodiment of the invention, a data management method comprises receiving a baseline version of a data object of a computer system using one of a plurality of storage devices of a data management system, storing the baseline version using the one of the storage devices after the reception of the baseline version, receiving a request using the one of the storage devices, wherein the request comprises a request to store a delta version of the baseline version, analyzing a capacity of the one of the storage devices, and storing the delta version using an other of the storage devices responsive to the analyzing.

[0011] Other embodiments are disclosed as is apparent from the following description.

DESCRIPTION OF THE DRAWINGS

[0012] Fig. 1 is a functional block diagram of an exemplary computer network according to one embodiment.

[0013] Fig. 2 is a functional block diagram of an exemplary data management system according to one embodiment.

[0014] Fig. 3 is a flow chart of an exemplary methodology for storing baseline and delta versions of a data object according to one embodiment.

[0015] Fig. 4 is a functional block diagram of an exemplary methodology for restoring the data object according to one embodiment.

DETAILED DESCRIPTION OF THE INVENTION

[0016] Referring to Fig. 1, a computer network 10 is shown according to one possible embodiment to depict exemplary embodiments described herein. Other configurations of computer network 10 are possible, including more, less,

and/or alternative components. As shown, the exemplary computer network 10 includes a network connection 12, a plurality of computer systems 14, and a data management system 16. At least some computer systems 14 may be referred to as protected computer systems and data management system 16 may be referred to as a resource domain in one embodiment. Data management system 16 may be configured to implement backup methods including disk archive where multiple versions of backup data objects may be stored on physical storage space (e.g., hard disk). As described further below, versions stored can be entire data objects or pieces of data objects. In one embodiment, an original version of a data object may be provided at an initial moment in time and may be referred to as a baseline or baseline version. Some subsequent versions of the data object provided at subsequent moments in time may only reflect changes with respect to the baseline version as opposed to comprising an entire version of the data object and may be referred to as delta versions.

[0017] Network connection 12 is configured to implement data transfer between coupled devices. In one arrangement, network connection 12 provides communications between computer systems 14 and data management system 16. Network connection 12 may also provide data communication intermediate internal devices of data management system 16 described further below in accordance with exemplary embodiments. The network connection 12 may comprise a local area network (LAN) and/or a storage area network (SAN) in possible configurations.

[0018] Computer systems 14 are configured to generate electronic data to be stored for subsequent retrieval and access (e.g., backup, restore, etc.). Exemplary computer systems 14 may comprise personal computers, work stations, servers, combinations of such devices, and/or other electrical devices capable of providing or accessing electronic data. In one embodiment, computer systems 14 comprise respective electronic file systems or groups of electronic file systems.

[0019] Data management system 16 is configured to store (protect) electronic data of computer systems 14. Data management system 16 may operate to perform exemplary operations, such as backup and restore, with

respect to electronic data of computer systems 14. Other operations may be provided.

[0020] As described below with respect to Fig. 2, data management system 16 may comprise resources configured to store a certain amount (e.g., bits or bytes) of electronic data. The data management system 16 accordingly may have a physical storage capacity corresponding to the physical size of the storage resources. In addition, data management system 16 may have a throughput or processing capacity corresponding to available data processing resources or additional quantifiable capacities in one embodiment.

[0021] Referring to Fig. 2, an exemplary configuration of a data management system 16 is shown. The depicted data management system 16 comprises storage control circuitry 20, a data storage system 22, and one or more client agent 24 coupled with network connection 12. Other arrangements of data management system 16 are possible.

[0022] In one embodiment, storage control circuitry 20 is implemented using a cell manager 26 comprising a server. An exemplary cell manager 26 comprises an interface 28, processing circuitry 30, memory 31, and a database 32.

[0023] Interface 28 is configured to display status information and receive user inputs in one arrangement. Interface 28 may be embodied as a graphical user interface (GUI) or other appropriate configuration.

[0024] Processing circuitry 30 is arranged to monitor and control operations of data management system 16. Processing circuitry 30 may be implemented using any appropriate configuration of hardware, software, and/or firmware to implement monitoring and control operations. For example, in the illustrated exemplary embodiment, processing circuitry 30 is implemented using a microprocessor. Memory 31 is configured to store processor-usable code or data for use by processing circuitry 30 in one embodiment.

[0025] Database 32 may be implemented in any appropriate medium and may be referred to as a tracking database. Database 32 is configured store configurations of the data management system 16, allocation of resources of system 16, relationship of the resources to computer systems 14, locations of stored data within system 16, and other desired data. As described further

below, database 32 stores location information regarding locations of baseline versions and delta versions of data objects within respective ones of a plurality of storage devices 34 of the data management system 16. Individual storage devices 34 may comprise an interface 35, an archive agent 36, and a physical storage space 38 in but one arrangement.

[0026] Interface 35 is configured to implement external communications of storage device 34 with respect to client agents 24, cell manager 26, other storage devices 34, and other external devices. Interface 35 may comprise a network interface card (NIC) in one embodiment although other configurations to implement communications are possible.

[0027] Archive agents 36 access data from appropriate client agents 24 and control storage of the data within the respective physical storage spaces 38. Archive agents 36 control reading and writing operations with respect to the respective physical storage spaces 38. Archive agents 36 may implement communications of data with respect to other storage devices 34, implement data compression and decompression operations, and receive and/or output data to computer systems 14.

[0028] Archive agents 36 include respective processing circuits 37 configured to control operations of the respective data storage devices 34. Processing circuitry 37 may be arranged using any appropriate configuration of hardware, software, and/or firmware (e.g., microprocessor and appropriate programming in the illustrated example) to implement monitoring and control operations. Memory 39 of the individual archive agents 36 are configured to store processor-usable code usable by the respective processing circuits 37. Archive agents 36 may comprise servers in one embodiment.

[0029] Archive agents 36 may additionally perform maintenance operations with respect to the data stored using the respective storage spaces 38. For example, in one embodiment, a baseline or a base version of a data object (e.g., word processing document, email, etc.) is stored using one of the physical storage spaces 38 at an initial moment in time. A change to the data object provides a delta version at a subsequent moment in time which includes the respective changes to the baseline version. Archive agents 36 may perform re-baseline maintenance operations wherein delta changes are applied to original

baselines of the data object, resulting in new baselines of the data objects at the moment of re-baselining. In addition, the delta version(s) may be combined during restoration operations to provide restored versions of the baselines. For example, the archive agents 36 communicate compressed data comprising the baselines and the deltas to the client agents 14 for the restoration wherein the delta(s) are applied to the respective baselines during restoration operations providing restored versions of the data object. Alternately, archive agents 36 may perform recombination operations wherein the delta changes are applied to the baseline version of a data object to provide a restored version of the data object which may be communicated to a client agent 24. Further, archive agents 36 may implement interleaving operations to communicate compressed data to respective client agents 14.

[0030] Physical storage spaces 38 may be implemented using any appropriate storage medium. Hard disk storage media may be utilized in at least one embodiment, although other storage media are possible (e.g., tape, semiconductor, holographic, etc.). Data storage system 22 comprises one or more of the physical storage spaces 38 in one illustrative arrangement. Data storage system 22 is configured to store data of a plurality of client protect computer systems 14. One exemplary embodiment provides a resource domain wherein a plurality of storage devices 34 are aggregated together and associated with one (or more) cell manager 26.

[0031] During configuration of data management system 16, client agents 24 may be associated with respective ones of computer systems 14. Client agents 24 may comprise programming (e.g., software) modules installed upon the respective computer systems 14. Other arrangements are possible.

[0032] In one embodiment, client agents 24 are configured to provide data from the respective computer systems 14 to the storage devices 34 for storage. Client agents 24 operate to extract data from the respective computer system 14, to compress the data (if desired), and to forward the data to be protected to a respective archive agent 36 for storage. Data communicated from a plurality of client agents 24 to one or more storage devices 34 may be interleaved (e.g., a large number of systems 14 are sending delta versions) wherein compressed data from a plurality of systems 14 is communicated to storage device(s) 34 in a

single sequential data transfer. During a restore, the archive agent 36 may resend the baseline and delta versions (potentially interleaved) to the respective client agents 14 in an appropriate order for the restore, or provide a complete restored version of the data object.

[0033] As described above, extracted and forwarded data may comprise original versions of data objects (baselines), and updated versions of the objects (deltas). Client agents 24 also operate to decompress received data from storage devices 34 (if the data is compressed) and to provide the data as live data for usage within the respective computer systems 14.

[0034] As mentioned previously, cell manager 26 may maintain a map of a distribution of the computer system 14 across the storage devices 34 using database 32 comprising a tracking database in one embodiment. As described further below, database 32 may store information regarding locations (e.g., identifying the storage devices 34) wherein delta versions and baseline versions are stored. The information may identify which storage devices 34 store baseline versions and respective delta versions for respective data objects. Cell manager 26 may be additionally configured to store information regarding associations of client agents 24 and computer systems 14 in tracking database 32. Accordingly, during exemplary storage operations, cell manager 26 maintains records in tracking database 32 identifying data objects and storage locations of data stored in the respective storage devices 34, including locations of baseline versions and delta versions of data objects.

[0035] In one embodiment, it may be desired to store all delta versions of a data object on the same storage device 34 which stores the respective baseline version of the data object. Such facilitates maintenance operations including, for example, re-baselining using the respective archive agent 36 of the storage device 34 or restoration. In addition, provision of the delta versions on the same storage device 34 which stores the baseline version minimizes degradation of LAN/SAN performance during maintenance processes (e.g., re-baselining) by providing internal data transfers within the respective storage device 34. Further, various pieces of a data object which are recombined for restore may be stored on different physical disk logical unit numbers (LUNs) if a common archive agent 36 is reading the data from the LUNs in one embodiment.

[0036] During data protection operations, resources of the data management system 16 including storage devices 34 are consumed. Consumed resources may include processing capacity (e.g., processing circuits 37) for processing operations with respect to protection of data, and storage capacity (e.g., physical storage space 38) for physically storing digital data. Over time, the amount of data to be protected of one or more computer system 14 typically increases. It may be possible to increase a capacity of physical storage space 38 (e.g., add additional disks) to accommodate the increased amount of data. Exemplary embodiments regarding scaling of storage resources according to one embodiment are described in a U.S. Patent Application, filed the same day as present application, entitled "Data Management Systems, Articles Of Manufacture, and Data Storage Methods," having client docket number 100204110-2, naming Stephen Gold, Harald Burose, and Sebastien Schikora as inventors, and the teachings of which are incorporated herein by reference.

[0037] However, some storage device configurations have maximum capacities wherein data may be properly protected. At some threshold, limitations may be imposed by an inability to efficiently add additional physical storage space 38, to efficiently process input/output requests, or to efficiently perform other processing in a timely manner within respective processing circuits 37, and/or other limitations. Without limits, the storage device may be a performance bottleneck for backup and restore or other operations.

[0038] According to at least one embodiment, data management system 16 is configured to continue data protection operations with respect to a storage device 34 even though a capacity of the respective storage device 34 is at a limit (e.g., the consumption of capacity of a storage device 34 is approaching or has reached a limit). In one embodiment, a storage device 34 may receive a request to store a delta version of an internally-stored baseline version, but have insufficient capacity (e.g., physical space, processing, etc.) to accommodate storage of the delta version. As described further below, data management system 16 may store the delta version at another storage device 34.

[0039] Referring to Fig. 3, an exemplary method implemented by processing circuitry 37 of a given storage device 34 with respect to storage of a baseline version and respective delta version(s) of a single data object is shown.

In the described exemplary embodiment, the physical storage space 38 of the storage device 34 comprising the processing circuitry 37 implementing the depicted method stores the baseline version of the data object, and this storage device 34 may be referred to as a primary storage device. Other storage devices 34 store delta versions of the baseline version and may be referred to as secondary or slave storage devices. Other methods are possible including more, less or alternative steps.

[0040] At a step S10, the processing circuitry receives a baseline version of a data object from a client agent.

[0041] At a step S12, the processing circuitry stores the baseline version using the associated physical storage space of the respective storage device. The processing circuitry may verify or otherwise indicate storage of the baseline version to the cell manager for tracking purposes using a tracking database.

[0042] At a step S14, the processing circuitry accesses a request to store a delta version corresponding to the stored baseline version of the data object.

[0043] At a step S16, the processing circuitry may initially attempt to store the delta version on the physical storage space of the primary storage device. The processing circuitry may monitor the capacity of the respective storage device. For example, the processing circuitry can obtain status information regarding processing and/or storage capacity of the storage device. Other methods may be used to provide capacity information.

[0044] If sufficient capacity exists (e.g., with respect to predefined limits), the processing circuitry may control the storage of the delta version using the physical storage space of the respective primary storage device at a step S18. The processing circuitry may then proceed to step S14 to process subsequently received delta versions.

[0045] Otherwise, if sufficient capacity does not exist, the processing circuitry proceeds in accordance with the described embodiment to a step S20 to output the delta version for storage using another storage device allocated to the same data protection solution as the primary storage device. The recipient storage device stores the delta version of the data object. In one embodiment, the processing circuitry selects the recipient storage device. In another

embodiment, the cell manager may dictate the appropriate recipient storage device. Other implementations are possible.

[0046] At a step S22, the processing circuitry is configured to output information to the cell manager for use in updating the tracking database to reflect the appropriate storage location of the respective delta version within another storage device.

[0047] Referring to Fig. 4, an exemplary method implemented by processing circuitry 37 of the primary storage device 34 with respect to restoration of a data object is shown. Other methods are possible including more, less or alternative steps.

[0048] At a step S30, the processing circuitry accesses a restore request with respect to a given data object. In one configuration, the request for restoration is initially sent by a client agent to the primary storage device which stores the baseline version for the given data object.

[0049] At a step S32, the processing circuitry may query the tracking database of the cell manager to obtain location information or data indicating the storage locations of baseline versions and delta version(s) of the respective data object responsive to a request for restoration of the data object. The tracking database may identify the request as requesting restoration of a data object which has data stored on at least one secondary storage device (e.g., one or more delta versions) as well as the primary storage device (e.g., baseline version and one or more delta versions). According to at least one embodiment, the processing circuitry of the primary storage device may initiate a distributed restore operation to provide an accurate point-in-time version of the data object to be restored and which includes data stored on one or more other storage devices.

[0050] At a step S34, the processing circuitry operates to receive the location data from the cell manager and which identifies the stored locations (e.g., locations of the secondary storage devices) which store data (data of delta versions) for the respective data object being restored.

[0051] At a step S36, the processing circuitry accesses the locally stored baseline version and delta version(s), if any, from the respective physical storage space of the primary storage device.

[0052] At a step S38, the processing circuitry accesses the remotely stored delta version(s) (i.e., delta version(s) not stored using the physical storage space of the primary storage device) using the location data obtained from the cell manager. In one embodiment, the processing circuitry of the primary storage device controls the respective primary storage device to operate as a proxy with respect to the secondary storage device(s) to extract the remotely stored delta version(s).

[0053] In one delta version data retrieval embodiment, the primary storage device sends one or more requests to secondary storage devices to provide specific delta versions. Responsive to receiving a request within a secondary storage device, the processing circuitry of the secondary storage device may retrieve the delta version and prepare the delta version for communication to the primary storage device as opposed to combining the delta version with other data during a typical restore operation. In some configurations, the stored delta version data may be compressed. In one embodiment, the archive agent of the secondary storage device may uncompress the delta version data before communicating the data to the primary storage device. In another embodiment, compressed data is communicated from the secondary storage device to the primary storage device.

[0054] At a step S40, the processing circuitry executes a recombination procedure to combine the data of the baseline version with the data of the accessed local and remote delta version(s) to provide a restored version of the data object.

[0055] At a step S42, the processing circuitry controls communication of the restored version of the data object to the appropriate client agent or other recipient.

[0056] Processing circuitry of the primary storage device may also verify the consistency of the data to be restored. In one embodiment, the processing circuitry of the primary storage device uses a Cyclic Redundancy Check (CRC) to verify the signature delta blocks, as well as an overall CRC to verify correct restoration of a complete restored data object. More specifically, the processing circuitry of the primary storage device may check the CRC data of the delta versions as they are read from the secondary storage devices, and a file-level

CRC may be generated in real time as the data object is being re-combined prior to restoration.

[0057] Delta versions may be communicated from secondary storage devices to the primary storage device using a plurality of different processes. A LAN-based process or SAN-based process may be used in illustrative examples.

[0058] In one exemplary LAN-based restore, a secondary storage device can extract requested delta versions from the associated physical storage space and provide the delta version data in a compressed state in one embodiment. For example, the data may have been extracted in a compressed state, or the processing circuitry of the secondary storage device may compress the data if it was uncompressed during extraction, or is otherwise provided in an uncompressed state. The compressed data of the delta version(s) may be communicated over a LAN or network connection to the primary storage device. The processing circuitry of the primary storage device may recombine delta versions read from the secondary storage devices with the baseline and delta versions stored using the primary storage device. The primary storage device may thereafter forward the restored data object as compressed data to a respective client agent. The recipient client agent may de-compress the data before writing it disk or other storage of the respective computer system. The client agent may provide the restored data in the same format as the original data object (e.g., file, raw volume for raw databases, or imported into an email application for single email objects).

[0059] In another embodiment, a SAN-based restore may be used for some types of data objects (e.g., excluding single emails which may be imported via an email application or raw data in one embodiment). For example, the physical storage space of a target computer system may be coupled with a SAN of the network connection. The primary and secondary storage devices may also be operatively coupled with the same SAN of the network connection. Processing circuitry of the secondary storage devices may operate to de-compress the delta version data prior to communication via SAN data movers to the primary storage device. Processing circuitry of the primary storage device may operate to uncompress locally stored baseline and delta version data, and combine the data with uncompressed data received from the secondary storage

devices. Following restoration, the primary storage device may communicate the restored version of the data object as uncompressed data to the appropriate client agent and computer system using SAN data movers in the described embodiment.

[0060] According to the illustrative SAN-based restore, the processing circuitry of the primary storage device may store restored data objects to a scratch area of its associated physical storage space (or other storage), and transfer portions (e.g., 100 MB) of the restored data objects over the SAN as subsequent portions are decompressed and/or recombined. Similar communications may be provided intermediate secondary storage devices and the primary storage device wherein the second storage devices may communicate un-compressed delta version data over the SAN to scratch space on the physical storage space of the primary storage device. This exemplary SAN configuration removes significant data transfer from the LAN or other connections, but provides communication of twice the amount of data compared with compressed communications and the archive agent of the primary storage device may bottleneck data transfer.

[0061] The client agent of the respective target computer system to receive the restored data object may reserve space on a physical array LUN for data objects to be restored using SAN-based restore. The respective client agent may operate as a resolve agent to decode a layout of volumes on primary disk LUNs to identify a list of physical blocks receiving the restored data transferred via the SAN data movers.

[0062] In another embodiment, one or more of the client agents may be configured to implement distribution of delta version data using a plurality of storage devices. In one implementation, the client agent may operate as a proxy to extract delta versions from the storage devices, and perform recombination operations of the baseline version and the delta versions. Data transfer intermediate the storage devices and the appropriate client agent may be performed using a LAN of the network connection inasmuch as the client agent may be performing the recombination operations and writing of the restored data to primary disk of the computer system and independent data transfer (e.g., SAN data movers) is not utilized in one example. Restoration operations may

slow performance of the LAN and client agent, and raw disk data objects may be more difficult to restore in embodiments wherein raw disk data objects are stored as whole LUNs permitting copying of such objects to file systems of the client agent.

[0063] At least some embodiments described herein may be implemented using processor-usable code or executable code stored within appropriate storage devices, or communicated via a network or using other transmission media, and configured to control appropriate processing circuitry. For example, processor-usable code may be provided via articles of manufacture, such as an appropriate processor-usable medium, or alternately embodied within a data signal (e.g., carrier wave, data packets, etc.) communicated via appropriate media, such as a communication network (e.g., the Internet and/or a private network) or other communication structure. Processor-usable media may include any media which can contain, store, or maintain programming for use by or in connection with an instruction execution system including processing circuitry in the exemplary embodiment. For example, exemplary processor-usable media may include any one of physical media such as electronic, magnetic, optical, electromagnetic, infrared, or semiconductor media. Some more specific examples of processor-usable media include, but are not limited to, a portable magnetic computer diskette such as a floppy diskette, or zip disk, hard drive, random access memory, read only memory, flash memory, erasable programmable read only memory, optical disk, or other configurations capable of storing programming, data, or other digital information.

[0064] According to at least some embodiments of the disclosure, solutions are provided for providing storage operations wherein a storage device of a data management system has reached a state wherein additional data may not be stored using the respective storage device. In one embodiment, an end user may scale the capacity of the data management system without having to reconfigure associations of individual storage devices to respective computer systems, and thereby avoid communication of a large amount of protected computer datasets intermediate plural storage devices. Additionally, performance degradation resulting from movement of significant amounts of data may be avoided.

[0065] The protection sought is not to be limited to the disclosed embodiments, which are given by way of example only, but instead is to be limited only by the scope of the appended claims.